

Abstract of the Disclosure

1
2
3 The invention provides a method and system for caching information
4 objects transmitted using a computer network. A cache engine determines directly
5 when and where to store those objects in a memory (such as RAM) and mass storage
6 (such as one or more disk drives), so as to optimally write those objects to mass storage
7 and later read them from mass storage, without having to maintain them persistently.
8 The cache engine actively allocates those objects to memory or to disk, determines
9 where on disk to store those objects, retrieves those objects in response to their network
10 identifiers (such as their URLs), and determines which objects to remove from the cache
11 so as to maintain sufficient operating space. The cache engine collects information to be
12 written to disk in write episodes, so as to maximize efficiency when writing information
13 to disk and so as to maximize efficiency when later reading that information from disk.
14 The cache engine performs write episodes so as to atomically commit changes to disk
15 during each write episode, so the cache engine does not fail in response to loss of power
16 or storage, or other intermediate failure of portions of the cache. The cache engine also
17 stores key system objects on each one of a plurality of disks, so as to maintain the cache
18 holographic in the sense that loss of any subset of the disks merely decreases the
19 amount of available cache. The cache engine also collects information to be deleted
20 from disk in delete episodes, so as to maximize efficiency when deleting information from
21 disk and so as to maximize efficiency when later writing to those areas having former
22 deleted information. The cache engine responds to the addition or deletion of disks as
23 the expansion or contraction of the amount of available cache.

BEST AVAILABLE COPY